The Entropy Game

Author: J Oliver Linton

Address: Pentlands, Keasdale Road, Milnthorpe, LA7 7LH

Email: jolinton@btinternet.com

Abstract: Using a simple computer model (which is fully described) the concept of entropy is discussed and the reasons why entropy always increases in systems that, on the face of it, appear to be symmetric with respect to time is clarified.

The second law of thermodynamics claims that in a sufficiently large system of an appropriate type, entropy always increases.

There are four problems with this statement:

- How large is 'sufficiently large'?
- What constitutes a system of the 'appropriate type'?
- What exactly is 'entropy'?
- To what extent does 'always' mean 'always'?

This article will help you clarify these issues.

A simple dynamic system

Imagine a grid of 64 by 64 squares, each of which can be either black or white. Rather like Conway's game of LIFE, at each 'generation', the board is updated according to a simple rule which is designed so that a small number of 'seeds' will gradually grow and grow. To get some idea of what the game does, figure 1 shows what the board looks like when a 5×5 square grows for 125 generations.

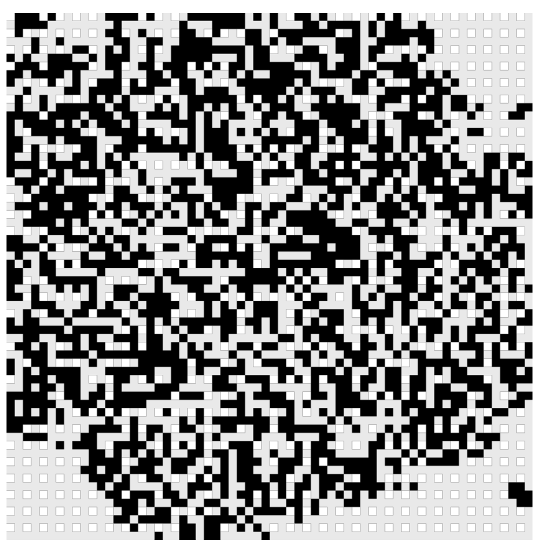

Fig 1: The Entropy Game after 125 generations

So what is the rule of the game? It is this.

The board is scanned from the top left hand corner in horizontal lines from top to bottom. For each cell in turn, the number of black cells in the adjoining 8 cells are counted. If this count is either 3 or 5, then the colour of the cell in question is reversed otherwise it is left alone. (This individual process is called an 'event'.) Since there are 4096 cells on the board, 4096 events occur every 'generation'. Now in Conway's game of LIFE, all the cells are scanned fist and then they are all updated all at once. The game of LIFE is *irreversible* (i.e. it cannot be run backwards) because there are, in principle, many different configurations that can give rise to a given pattern. In the Entropy Game, however, the cells are updated one by one immediately. The advantage of this system is that any operation can instantly be undone. All you have to do is apply the same algorithm but scan the cells in the reverse order. It follows that if you start with the configuration illustrated in figure 1 and run the game backwards, it will all unwind and after precisely 125 generations you will return to a 5×5 square!

Similarities with Newtonian Dynamics

It is a well known feature of Newton's laws that they are completely reversible. In principle, therefore, if you were to instantaneously reverse the velocities of all the balls on a snooker table just after the break, they would all collide with each other in just the right way to reconstitute the triangle and expel the cue ball. (This does, of course assume that the table is perfectly smooth and that the balls are perfectly elastic. This is an important point to which we shall return.) This is why the Entropy Game is a good model of Newtonian Dynamics. It is deterministic and reversible. It really doesn't matter what the rules of the game are provided that they produce patterns which develop over time and which are sensitive to the original conditions.

The best system in which to see Newtonian Dynamics at work is the classic one of gas molecules in a box. There is no friction in empty space and gas molecules are indeed perfectly elastic (provided the collisions are not to violent). So if you were to release some gas in the corner of a box and then instantly reverse their velocities exactly one second later, you would expect all the gas to collect back in the corner again, wouldn't you?

Well, we all know that this wouldn't happen in reality, but why not? Is it because it is impossible in practice to reverse the velocities exactly? Is it because there is some friction in space after all? Is it because even gas molecules are not perfectly elastic? Or is it simply because the second law of thermodynamics forbids it in the same way that the first law forbids perpetual motion?

In fact, none of these solutions to the riddle are correct in spite of the fact that many millions of words have been written in defence of one or other (especially the first).

I believe that the correct answer is that molecules do not obey Newton's laws exactly, they obey the laws of quantum mechanics and the crucial thing about QM is that particles such as gas molecules bo not possess a precise position and velocity (or to be more accurate, position and momentum). As is well known the product of the uncertainty in each of these quantities cannot be less than $h/2\pi$ (where h is Plank's constant = 6.6 x 10^{34} Js) The effect of this is to introduce a tiny element of randomness into every collision. Now the number of collisions which the molecules of a gas in a bottle make in one second is staggeringly large (It is of the order of 10^{35}) and if even one of these collisions goes wrong, the errors will accumulate exponentially and soon all correlation between the molecules will be irretrievably lost. So, unlike the Entropy Game – which is strictly reversible and deterministic – the 'entropy' of gas molecules in a box (and real snooker balls on a real snooker table) always increases in practice. But how do we define and calculate 'entropy'?

Calculating 'entropy'

'Entropy' is sometimes described a s a measure of the disorder of a system. If a system is disordered, it takes a lot of words to describe it (Imagine having to write down the exact position of every toy in a playroom after the children have departed) whereas if the system is highly ordered, it will take far fewer words to described it ('All the toys are in the toy cupboard'). Likewise, a 5×5 square in the middle of the 64×64 cells of the Entropy Game board is a lot easier to describe than the pattern illustrated in figure 1.

In the Entropy Game we shall define the 'entropy' of any pattern as the minimum number of binary digits ('bits') which are needed to describe the pattern completely. But in order to be able to distinguish between highly ordered patterns and highly disordered ones, we are allowed to use the best 'compression algorithm' that we can devise. (The one I describe here is far from being the best, but it will do.)

The simplest way to describe any 64×64 pattern is in the form of a binary number with 4096 bits, 0 representing a white cell and 1 a black. I shan't write out the whole number which corresponds to the pattern in figure 1 – it would waste too much paper – but it begins and ends like this:

111010 ….. 0000000

If, on the other hand, the pattern has a bit of order in it there are shorter ways of describing it. For example, suppose a row of 64 cells looks like this:

0000000001111111111111100001111111111111111111111111111111110000

we could encode each row as follows: first we note that there are 5 groups. There are 9 zero's followed by 14 one's, 4 zero's 33 one's and the rest are zero's so the pattern is largely described by the four numbers 9, 14, 4, 33. But we need to prefix these numbers with a bit more information. Firstly the number of numbers in the group (4 in this case) and also the identity of the first bit (0 in this case). So the complete code for the line (in decimal notation) for the line is:

4, 0, 9, 14, 4, 33

Now a bit of thought will convince you that it is sufficient to allow 3 binary digits for the first number (4), one bit for the second and 5 bits each for the rest so the following binary number will do the trick:

010 0 01001 01110 00010 10001

(I have put spaces in between the groups to help you sort them out but these are unnecessary) It is obvious that this contains far fewer bits (24 in fact) than the original which has 64 bits.

There is just one problem. If a line has more than 7 groups in it, it will need more than 64 bits to encode it and the first number will need more than 3 bits to define it. We therefore limit the number of groups to a maximum of 6 and if there are more than this, revert to specifying the line in full prefixing the line with a special 4 bit code 1111. This means that a complicated line needs 68 bits to describe it.

Once you have calculated the code for each line, you just string them all together to get the code for the whole board and the 'entropy' of the pattern is simply the total number of bits in the whole code!

Lets calculate the entropy of the 5 x 5 square: The first 30 lines all have the code 000 0 (there is only one group - 'the rest' - and they are all zero) i.e 4 bits. The next 5 lines have the code 010 0 11110 00101 (or 2, 0, 30, 5 in decimal) i.e. 14 bits ,and the last 29 lines need 4 bits each like the first ones. The total number of bits needed is 30 x 4 + 5 x 14 + 29 x 4 = 306.

The entropy of a completely random board maxes out at 64 x 68 = 4352

Entropy in the real world

This is all very well, but does this method of calculating entropy have anything to do with calculating entropy in the real world? Yes it does. In textbooks on thermodynamics you will find entropy S defined as:

$$S = k \log V$$

k is just a constant (Boltzman's constant) and V is the number of different ways in which the system can be arranged and still 'look the same'. Don't worry too much about what 'looks the same' means in the context of a thermodynamic stsyem; it has a pretty obvious meaning in the context of the Entropy Game. All completely random patterns 'look the same' and have the same 'entropy' (i.e. 4352). Any pattern with some recognisable order in it will not 'look the same' and will have a lower 'entropy'. (Owing to the particular way I have defined my compression algorithm, a vertical bar will not have the same 'entropy' as a horizontal one but it won't be that far different and the differences can safely be ignored.)

In counting the number of bits needed to describe the pattern, I have effectively taken the logarithm (to base

2) of the number of patterns which can be described using a certain number of bits and which therefore, according to my rules, 'look the same' and this is precisely what entropy is - the logarithm of the number of possible patterns which 'look the same' according to the rules of the game!

How large is 'sufficiently large'?

At the beginning of this essay I posed four questions: The first is how large must a system be to display the features of the second law? The answer to this is simple. It must have so many possible configurations that it is impossible to exhaust them within a reasonable time-scale.

Let us calculate how long it would take to run through all the possible configurations of a 64×64 board. Since each cell can be in one of 2 states, there are just 2^{4096} different patterns. Now it is difficult to conceive just how staggeringly large this number is. It is easy to make the mistake of thinking that 2^{4096} can't be very different from 4096^2, after all 2^3 is not very different from 3^2 is it? But 4096^2 less than 17 million whereas 2^{4096} is a number which has 1233 decimal digits. Even the fastest computer in the world working for the age of the universe could not possibly make the tiniest dent in a number this large!

Mathematicians like to use a device called 'configuration space' to explore the properties of dynamic systems. If a system consists of N independent quantities each of which can take on one of p possible values, then any particular state of the whole system can be represented by a single point in a space which has N dimension and p 'ticks' on each dimension. For example, consider a multicolour LED whose three components (Red, Green and Blue) can be either on or off. In this system $N = 3$ (there are three independent quantities, R, G and B) and $p = 2$ (each of which can be on or off). It is easy to see that there are $p^N = 2^3 = 8$ possible states in this system and they can be conveniently visualised as the corners of a cube. (see figure 2).

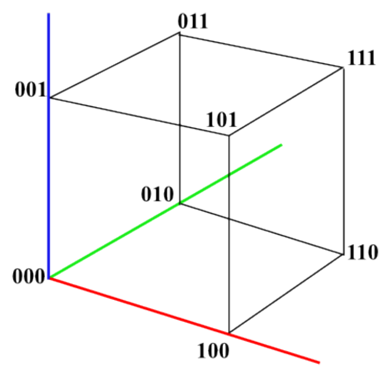

Fig 2: Configuration space of a 3-colour LED

Now we shall assign some 'entropy' values to the 8 possible states. First we note that the 8 states divide into 4 groups which 'look the same'. They are [000], [001, 010,100], [011. 101, 110] and [111]. (Note the numbers of states in each group namely 1, 3, 3 and 1) The states 000 and 111 are obviously less 'disordered' than the states others but why is this? It is because, as a group, they require fewer words to describe then (e.g. 'All on' as opposed to 'all off except the green') Nor is it an accident that there are fewer of them. In fact we can define the 'entropy' of any state as the logarithm of the number of states that share the same overall description. So the 'entropy' of the state 000 is log(1) = 0, the entropy of states like 001 is log(3) = 0.48, states like 101 have the same 'entropy and the 'entropy' of 111 is again 0. (I have used logs to base 10 for convenience but any base will do – the *ratio* of the figures will be the same.)

Now imagine taking a random walk round the edges of the cube. At each corner you can move in any of three directions i.e. there are 24 possible choices to make. (e.g. I am standing at 101 and I choose to move in the red direction) Of these 24 possibilities, 6 will result in an *increase* in 'entropy' (the moves away from 000 and 111); 6 will involve a *decrease* in entropy (the moves towards 000 and 111) and the rest (12) will result in 'entropy' staying the same. So far there is no evidence of any bias in the figure which would suggest that anything like the Second Law is at work, but this is partly because the system we are considering is too small.

So lets consider 4-colour LED. Now there are 16 states arranged at the corners of a hypercube. The states fall into 5 groups with 0, 1, 2, 3 or 4 colours lit and the numbers in each of the 5 groups are 1, 4, 6, 4 and 1 with entropies 0, 0.6, 0.78, 0.6 and 0 respectively. (You will not have failed to notice that the sizes of the groups are simply the well-known binomial coefficients of $(x + y)^n$.)

Now lets make the big conceptual leap to a binary system with N dimensions (like the Entropy Game where $N = 4096$ and there are 2^N states). We can list all the states by groups in order and plot the 'entropy' of each group along a line. What we will get is something like the curve[1] in figure 3:

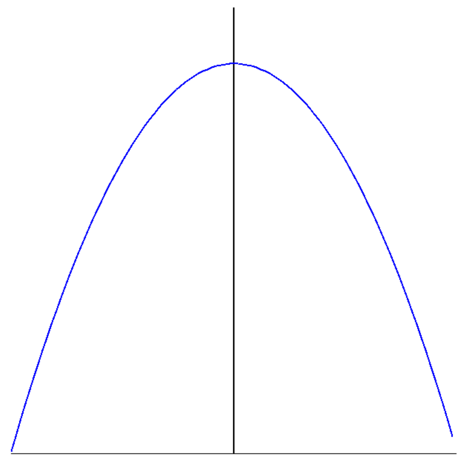

Fig 3: 'Entropy' values for a large binary system

If you were plonked down at random in one of the states of this system (lets say the k^{th} group), the chances are that the 'entropy' of the state you were in would be high but not maximal. Now we need to work out which way 'entropy' is likely to go if we take a random step in any of the N possible direction we can take.

It turns out that you are approximately $\left(\frac{n-k}{k}\right)^2$ times[2] more likely to move in a direction in which 'entropy' increases. So, for example, in the Entropy Game where $N = 4096$, if you are in, say the 1000^{th} group, you would be nearly 10 times more likely to move in the direction of increasing 'entropy'

But this raises an important issue. The Entropy Game is not random. It is all very well proving that in a random game, 'entropy' is almost bound to increase. All we are really saying is that in random systems probable outcomes are more likely than improbable ones. And who needs a mathematician to tell us that? The question is – why do deterministic systems like the Entropy Game and Newtonian dynamics obey the second law? What is special about them?

What constitutes a system of 'appropriate type'?

Well the first thing to say is that the Entropy Game does *not* obey the second law of thermodynamics. The pattern illustrated in figure 1 is not random. As I said earlier, it was generated by running a 5×5 square through 125 generations and if we were to run it though the game in reverse it would wind itself back to the start and its 'entropy' would decrease.

But if you were to change one single cell from black to white, there would be a totally different outcome. At first the entropy would decrease because the effects of single change are restricted to the 8 nearest neighbours of each cell on each generation. But these effects multiply and spread out – at the 'speed of sound' as it were – until the whole board is affected and 'entropy' increases again. Figure 4 shows how the calculated 'entropy' developed over 150 subsequent generations when a single cell in the corner of the board was changed.

1 This curve is the logarithm of the normal distribution I.e. $\log(\exp(-x^2))$

2 The ration of the k+1^{th} term of the binomial expansion to the k-1^{th} term is $(N-k)(N-k+1)/k(k+1)$

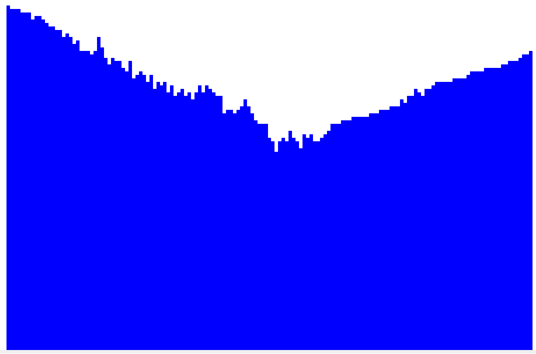

Fig 4: Calculated 'entropy' of a pattern with a single error

This suggests that there are two qualities which a system must possess if it is to exhibit the behaviour associated with the second law: firstly it must be extremely sensitive to small changes in its initial conditions and secondly, that it must contain a tiny bit of randomness. Take the snooker break that we mentioned earlier. Because of friction between the ball and the table, the reversed balls would not collide with each other in exactly the same (reversed) order and the eventual outcome will be totally different. The same is true of gas molecules in a box. The uncertainty in the position of a molecule when it strikes another molecule cannot be less than the Plank length of 10^{-35} m. Now since the average separation of molecules in a gas at atmospheric pressure is typically 10 times the radius of a molecule, when a molecule collides a second time it will typically have an uncertainty in its position of 10 times this and after 25 collisions, the uncertainty in its position will be 10^{-10} m (i.e. comparable to the radius of a molecule) and it will probably miss its target and collide with a completely different one. Now gas molecules make 10^{12} collision *every second* so we are fully justified is regarding the motion of gas molecules as being completely random.

There is, of course, one other condition which must be met. If entropy is going to increase, it must start in a configuration in which entropy is less than maximum. We are fortunate to live in a region of the universe whose entropy is staggeringly low and there is a long way to go before we are in any danger of running out of 'entropy space'!

Conclusions

Notwithstanding the huge number of words which have been wasted on explaining how entropy can always increase in a system which appears to be symmetrical with respect to time, I believe that the answer to the riddle of the relation between the Second Law of Thermodynamics and the 'Arrow of Time' is simply that the uncertainty implicit in Quantum Theory together with the extreme sensitivity to initial condition displayed by a large number of dynamical systems ensures that such systems do not behave symmetrically with respect to time. Although 'entropy' is a difficult concept to pin down and there are many different ways of defining it, at the end of the day, the precise method of calculating it turns out not to matter very much. Provided the system under study is 'sufficiently large' and is of the 'appropriate type' entropy, however it is defined always ends up increasing.

To misquote a famous politician, in practice 'always' means 'always'.

A computer model which runs on Windows is available on the author's website:

http://www.jolinton.co.uk